

REMARKS

Claims 1-20 are pending in the application. As indicated on the cover sheet for the Request for Continued Examination, the prior amendments filed, September 22, 2006, i.e., those presented after final rejection, should NOT be entered. The current amendments mirror the previously presented amendments but include some additional clarifications and limitations in independent claims 1 and 11, but otherwise do not change the scope of the prior amendments. (The current amendments also reduce the number of independent claims in comparison to the claim amendments presented after final rejection.)

Claims 1-20 stand rejected under 35 U.S.C., §112, second paragraph as being indefinite, and also stand rejected under 35 U.S.C. §102(e) as being anticipated by Aziz (U.S. Patent No. 6,597,956).

With this response, claims 1-4, 8-14, 16, and 18-20 have been amended to address the Examiner's §112, second paragraph concerns, and claim 6 was similarly amended to provide antecedent basis for "the plurality of computer processors."

Aziz discloses an extensible computing system (title) which can be used to create a Virtual Server Farm (VSF) from a wide scale grid of computing entities that contains tens of thousands (or more) of computing elements (see Abstract; see also col. 6, ll. 33-37). The system includes a "local computing grid" with "computing elements CPU1, CPU2,...CPUn" and storage area network (SAN) switches that couple the computing elements to "disks DISK1, DISK2,...DISKn" (col. 6, lines 35-44). The system also includes virtual LAN (VLAN) switches that interconnect the computing elements and that are coupled to the Internet (col. 6, lines 47-50). The system also includes a control plane, which is preferably implemented in a hierarchical master-slave manner and which is "not directly accessible through any of the computing elements in the computing grid," but rather "is implemented on a completely independent set of computing elements assigned for supervisory purposes," (col. 5, lines 44-45).

To construct a VSF, the control plane "assigns or allocates" CPUs to function as load balancers/firewalls, web servers, or other entities in the VSF (col. 10, lines 40-42). To configure a VSF, the control plane sends control commands to the various VLAN and SAN switches (col.

1, lines 19-21) through special control ports or interfaces of the networking and storage elements (see col. 5, ll. 49-51 and col. 14, ll. 18-28).

Once a VSF is configured, the computing entities of the VSF may directly communicate with one another and may communicate with external entities, e.g., on the Internet, or disk storage. This is shown in Figures 2 and 5 and corresponding text. For example, Aziz discloses that:

[T]he computer can access the storage device on the Fibre-Channel SAN just as it would access a locally attached SCSI disk. Therefore, software such as boot-up software simply boots off the disk device on the SAN just as it would boot off a locally attached SCSI disk (col. 12, lines 43-49, emphasis added).

The control plane is used to configure a dynamic firewall scheme (see, e.g., col. 5, line 28-31.). Among other things, the control plane is used to configure the SAN and VLAN switches to configure virtual LANs and SAN zones (see col. 6, l. 66 – col. 7, l. 8; see also col. 8, ll. 35-8; see also col. 8, l. 67 – col. 9, l. 1.) The control plane is not a conduit for storage or network messages to the disks or the external Internet, nor is there a need or motivation for such suggested in Aziz's firewalls scheme.

In contrast, the amended claims recite a platform and method in which messages from computer processors to the external communication network and to the external storage network are received and processed by the least one control node and in which the control node transmits processed forms of the messages to the external communication network and to the external storage network. Support for the claims exists in the specification. Pages 4-5 (and throughout the specification) provide high level description of the architecture. Pages 22-25 provide support for the processing external network I/O messages by the control node, and pages 32-40 provide support for the processing of storage network messages by the control node.

Aziz and the recited claims share a high level similarity in that both allow software (not manual) configuration of VSFs or processing area networks. After that, however, the similarity ends. Aziz is a form of grid computing system intended for use on extremely large collections of computers. Aziz for example refers to widely distributed systems that may contain tens of thousands of processors or more (see, e.g., col.6, line 35). Virtual LAN technology and SAN

zones are used to allocate computers into VSFs and to ensure security. The control plane is used to configure the VLANs and SAN zones , but it is not involved in the subsequent and normal (i.e., not configuration) processing of messages to network I/O or storage.

In contrast the claims are directed to a high performance configurable platform of processors. It includes an internal communication network connecting the various processors and a control node and configuration logic to configure the processors into processing area networks with specified network topologies and virtual storage spaces. The control node is connected to the internal network and also connected to and communicating with the outside world, i.e., the external communication network (e.g., network I/O, internet, etc) and external storage network (e.g., external SANs). It is the conduit for messages between the processors and the outside world. This claim language connotes and is consistent with the high performance platforms described in the specification, which generally fit in one chassis (or perhaps a very small number of chasses). People skilled in the art would not confuse the high performance platforms described or claimed with a grid computing system like Aziz.

The dependent claims contain many further distinctions over Aziz as well. For example, claim 8 recites, in part:

wherein the at least one control node receives, via the internal communication network, storage messages from said corresponding set of computer processors, and wherein the at least one control node includes logic to extract an address from a received storage message, to identify the defined corresponding address in the external storage space, and to provide messages on the external storage network corresponding to the received storage messages and having the corresponding address.

That is, the claim clearly recites a platform in which storage messages are received by the control node, and the control node actively operates on such.

Aziz does not disclose or suggest anything like the recited control node. Aziz's control plane does not receive storage messages from computer processors. Nor does Aziz's control plane include logic to extract an address from a received storage message, to identify the address in the external storage space, and then to provide messages corresponding to the received storage messages on the external storage network.

In rejecting claim 8, the Examiner cited sections of Aziz that simply disclose that the control plane can issue “control commands” to the VLAN and SAN switches via control ports (col. 10, lines 16-21), and that a communication interface can be used to provide “two-way data communication” between a computer system and a local network (col. 26, line 34-col. 27, line 6). However, the cited sections do not disclose or suggest that the control plane receives storage messages, extracts addresses from them, and then provides corresponding storage messages on the external storage network, as recited in claim 8 as amended. Thus, Aziz does not teach or suggest claim 8.

Claim 9 includes limitations from claim 8 and should be allowable for the above reasons. It also recites further novel features, namely:

wherein the at least one control node includes logic to buffer data corresponding to write messages received from a computer processor of said corresponding set of computer processors and to provide the buffered data in the corresponding message provided to the external storage network.

Aziz does not disclose or suggest a control node that includes logic to buffer data corresponding to write messages received from a processor, and then to provide the buffered data to an external storage network, as recited in claim 9. As discussed above, Aziz discloses configuring CPUs to directly access SAN disks “just as it would access a locally attached SCSI disk” (col. 12, lines 35-46). In rejecting claim 9, the Examiner cited a section of Aziz that describes ways to establish direct communication between the CPUs and SAN devices, e.g., with a SCSI-to-Fibre channel bridging device (col. 12, line 35-col. 13, line 16). This is the exact opposite of the claim. Neither the cited section, nor any other section of Aziz, discloses or suggest that the control plane includes logic to buffer data corresponding to write messages received from a processor, and then to provide the buffered data to an external storage network, as recited in claim 9.

Claim 10 includes limitations from claim 8 and should be allowable for the above reasons. It also recites further novel features, namely:

wherein the at least one control node receives storage messages from the external storage network, and wherein the at least one control node includes logic to identify a corresponding computer processor or control

node that the received message is responsive to, and to provide a corresponding message to the identified computer processor or control node.

Aziz does not disclose or suggest a control node that receives storage messages from an external storage network, includes logic to identify a processor or control node to which the message is responsive, and then provides a corresponding message to the processor or control node. As discussed above, Aziz discloses configuring CPUs and SAN devices to communicate directly with each other. In rejecting claim 10, the Examiner cited sections of Aziz that simply disclose that the control plane can issue “control commands” to the VLAN and SAN switches via control ports (col. 10, lines 16-21), and that a communication interface can be used to provide “two-way data communication” between a computer system and a local network (col. 26, line 34-col. 27, line 6, and Fig. 19). The cited sections, nor any other sections of Aziz, disclose or suggest that the control plane receives storage messages, includes logic to identify a processor or control node to which the message is responsive, and to provide a corresponding storage messages to the identified processor or control node, as recited in claim 10 as amended.

Applicants submit that claims 18, 19, and 20 include similar limitations to claims 8, 9, and 10, respectively, and should be similarly allowable for the above reasons.

Claim 5 recites a platform in which a control node communicates with computer processors on a particular kind of internal network, and is in the path between the processors, an external communication network, and an external storage network. Specifically, claim 5 recites “wherein the internal communication network is a point to point switch fabric.”

Thus, claim 5 distinguishes between the computer processors, the control node, and the internal communication network, i.e., point to point switch fabric, through which the control node communicates with the computer processors. Aziz discloses a “wide scale computing fabric (“computing grid”),” which the Examiner appears to equate with the “point to point switch fabric” of claim 5. However, the “wide scale computing fabric” or “computing grid” disclosed by Aziz is not even an internal communication network, let alone a point to point switch fabric, as is recited in claim 5. Instead, as Aziz discloses:

The computing grid comprises a large plurality of computing elements that are coupled to one or more VLAN switches and to one or more storage area network (SAN) switches. A plurality of storage devices are coupled to the SAN switches and may be selectively coupled to one or more of the computing elements through appropriate switching logic and commands (col. 3, lines 25-28, emphasis added).

Thus, when Aziz uses the terms “computing fabric” and “computing grid,” he means the computer processors themselves. Aziz does not disclose or suggest a point to point switch fabric through which the processors communicate with a control node, as recited in claim 5.

Applicants submit that as claim 15 includes similar limitations to claim 5, it should be similarly allowable for the above reasons.

In view of the above amendment, Applicants believe the pending application is in condition for allowance, and respectfully request the Examiner to allow the claims to issue.

Applicants hereby request that the period for responding to the outstanding Office Action be extended for two months’ time to maintain the pendency of the above-identified case. The Commissioner is hereby authorized to charge the required fee for filing the request for extension of time to our Deposit Account No. 08-0219. No other fees are believed to be due at this time. However, please charge any additional fees which may be due, or credit any overpayments, to Deposit Account No. 08-0219.

Respectfully submitted,

Dated: January 8, 2007



Peter M. Dichiara
Registration No.: 38,005
Attorney for Applicant(s)

Wilmer Cutler Pickering Hale and Dorr LLP
60 State Street
Boston, Massachusetts 02109
(617) 526-6000 (telephone)
(617) 526-5000 (facsimile)